

REMARKS

Claims 1-26 are pending in the present application.

Claims 1-3, 9-12, 18-20, and 26 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Mehrota et al. (U.S. Patent No. 6,145,054) (hereinafter " Mehrota ") in view of Gruner et al. (U.S. Patent Publication No. 2003/0154346) (hereinafter " Gruner"). Applicant respectfully traverses this rejection.

Claims 4-5, 13-14, and 21-22 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Mehrota in view of Gruner, and in further view of Flynn (U.S. Patent No. 6,345,335) (hereinafter "Flynn"). Applicant respectfully traverses this rejection.

Claims 6-8, 15-17, and 23-25 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Mehrota in view of Gruner, and in further view of Rowlands (U.S. Patent Publication No. 2003/0217236) (hereinafter "Rowlands"). Applicant respectfully traverses this rejection.

Applicant's claim 1 recites a cache memory subsystem comprising in pertinent part,

“a scheduler configured to schedule reads and writes of information to said cache storage using a fixed latency pipeline;
wherein in response to scheduling a read request to said cache storage,
said scheduler is further configured to cause an associated write to said cache storage to occur a fixed number of cycles after said scheduling a read request”

The Examiner acknowledges Mehrota is only relied upon to teach “a cache storage configured to store a plurality of cache lines of data.” The Examiner further asserts Gruner teaches the remaining limitations of claim 1. However, the Examiner asserts Gruner teaches the use of a pipeline “as described by Mehrota above.” Applicant respectfully disagrees with the Examiner's assertion that Mehrota teaches a pipeline as recited in Applicant's claim 1. Applicant notes that just because Mehrota mentions the

word “pipelined” in the context of a cache hierarchy, does not mean Mehrota teaches the limitation in Applicants claim 1. Specifically, Mehrota teaches at col. 2, line 65 through col. 3, line 11

“In a pipelined hierarchical cache system that generates multiple cache accesses per clock cycle, coordinating data traffic between the different cache levels is problematic. For example, when a first access to a given cache line results in a miss, the access is sent on to be serviced by a higher cache level or main memory. When the first access is completed, the cache line becomes valid. In typical cache organizations, after the cache line becomes valid, it is forwarded to a lower cache level or device that generated the original access. The cache line fill operation needs to be synchronized with the return data, but the lower level cache executing the line fill operation cannot predict when the required data will be returned.” (Emphasis added)

From the foregoing, Mehrota is merely identifying issues with a typical pipelined cache system. However, Applicant submits Mehrota does not teach or disclose anything about the nature of the pipelined cache hierarchy. In fact, Mehrota is teaching that one cannot predict when the data will be returned. This is clearly in contrast to having a fixed latency R/W pipeline.

The Examiner asserts Gruner “helps disclose a scheduler ... using a pipeline (as described by Mehrota above) with a fixed latency ...” Applicant respectfully disagrees with the Examiner’s characterization of Gruner. More particularly, Applicant is not sure what the Examiner’s recitation in Gruner is supposed to teach, but it clearly does not teach the limitations in Applicant’s claim 1. Specifically Gruner teaches at paragraphs [0117], [0118] and [0131]

“Although not shown, cache 80 includes receive data buffers, in addition to the request queues shown in stage 370. The receive data buffers hold data passed from cache 52 for use in requested memory operations, such as stores. In one embodiment, cache 80 does not contain the receive data buffers for data received from data ring 20 along with Fill requests, since Fill requests are serviced with the highest priority. Cache 80 includes a scheduler for assigning priority to the above-described memory requests. In stage 370, the scheduler begins the prioritization process by selecting requests that originate from snoop queue 390 and each of compute engines 50, 86, 88, and 90, if any exist. For snoop request queue 390, the scheduler selects the first request with a Validity field showing the request is valid.

In one embodiment, the scheduler also selects an entry before it remains in queue 390 for a predetermined period of time.

[0118] For each compute engine, the scheduler gives first tier instruction cache requests (FTI) priority over first tier data cache requests (FTD). In each data cache request queue (382, 384, 386, and 388), the scheduler assigns priority to memory requests based on predetermined criteria. In one embodiment, the predetermined criteria are programmable. A user can elect to have cache 80 assign priority based on a request's Opcode field or the age of the request. The scheduler employs the above-described descriptors to make these priority determinations.”

[0131] If the requested operation is a store, cache 80 performs a read-modify-write operation. Cache 80 supplies the addressed cache line to store buffer 393 in stage 380. Cache 80 modifies the store buffer bytes addressed by the first tier memory request. Cache 80 then forwards the contents of the store buffer to data array 396. Cache 80 makes this transfer once cache 80 has an idle cycle or a predetermined period of time elapses. For stores, no data is returned to first tier data cache 52. (Emphasis added)

From the foregoing, it appears Gruner is merely teaching a scheduler that prioritizes memory requests. Applicant cannot find any relevance to the limitations in Applicant’s claim 1. Clearly, in paragraph [0131] Gruner is teaching that the request is a store request, and not a read request, and it is not scheduling an associated write of data in response to scheduling a read request. Further, there is no mention of fixed latency or an associated write occurring some fixed number of cycles after the scheduling of the read.

Thus, Applicant submits neither Mehrota nor Gruner, taken either singly or in combination, teach or suggest the combination of features recited in claim. Specifically, neither Mehrota nor Gruner teach or suggest “a scheduler configured to schedule reads and writes of information to said cache storage using a fixed latency pipeline” or “wherein in response to scheduling a read request to said cache storage, said scheduler is further configured to cause an associated write to said cache storage to occur a fixed number of cycles after said scheduling a read request,” as recited in Applicant’s claim 1.

In regard to the Examiner's rejection of claim 2, Applicant respectfully disagrees with the Examiner's characterization of both Mehrota and Gruner and the Examiner's allegation of the combined teaching therein. Specifically, Applicant submits the combined teachings of Mehrota and Gruner merely teach conventional cache system victim data queue and filling of that data from one cache to another. Mehrota teaches servicing the cache miss with a line of data, and Gruner teaches using a victim queue, which teaches completely away from Applicant's invention. (*See* Applicant's background section, paragraph [0006].

Neither Flynn nor Rowlands were relied upon to teach the limitations recited in Applicant's claim 1. In addition, Applicant submits that neither Flynn nor Rowlands teach or suggest the features recited in Applicant's claim 1.

Accordingly, Applicant submits claim 1, along with its dependent claims, patentably distinguishes over Mehrota in view of Gruner for the reasons given above.

Applicant's independent claims 10 and 18 recite features that are similar to the features recited in claim 1. Thus, for at least the reasons given above in the discussion of claim 1, Applicant submits claims 10 and 18, along with their respective dependent claims, patentably distinguish over Mehrota in view of Gruner.

CONCLUSION

Applicant submits the application is in condition for allowance, and an early notice to that effect is requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5500-91000/SJC.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Stephen J. Curran", is written over a horizontal line.

Stephen J. Curran
Reg. No. 50,664
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: December 12, 2006